

CausalImpact: A new open-source package for estimating causal effects in time series

Posted: Wednesday, September 10, 2014



423



398



250

How can we measure the number of additional clicks or sales that an [AdWords](#) campaign generated? How can we estimate the impact of a new feature on app downloads? How do we compare the effectiveness of publicity across countries?

In principle, all of these questions can be answered through [causal inference](#).

In practice, estimating a causal effect accurately is hard, especially when a randomised experiment is not available. One approach we've been developing at Google is based on [Bayesian structural time-series models](#). We use these models to construct a synthetic control — what would have happened to our outcome metric in the absence of the intervention. This approach makes it possible to estimate the causal effect that can be attributed to the intervention, as well as its evolution over time.

We've been testing and applying structural time-series models for some time at Google. For example, we've used them to better understand the effectiveness of advertising campaigns and work out their return on investment. We've also applied the models to settings where a randomised experiment was available, to check how similar our effect estimates would have been without an experimental control.

Today, we're excited to announce the release of [CausalImpact](#), an open-source [R](#) package that makes causal analyses simple and fast. With its release, all of our advertisers and users will be able to use the same powerful methods for estimating causal effects that we've been using ourselves.

Our main motivation behind creating the package has been to find a better way of measuring the impact of ad campaigns on outcomes. However, the [CausalImpact](#) package could be used for many other applications involving causal inference. Examples include problems found in economics, epidemiology, or the political and social sciences.

How the package works

The [CausalImpact](#) R package implements a Bayesian approach to estimating the causal effect of a designed intervention on a time series. Given a response time series (e.g., clicks) and a set of control time series (e.g., clicks in non-affected markets, clicks on other sites, or [Google Trends](#) data), the package constructs a Bayesian structural time-series model with a built-in spike-and-slab prior for automatic variable selection. This model is then used to predict the counterfactual, i.e., how the response metric would have evolved after the intervention if the intervention had not occurred.

As with all methods in causal inference, valid conclusions require us to check for any given situation whether key model assumptions are fulfilled. In the case of [CausalImpact](#), we are looking for a set of control time series which are predictive of the outcome time series in the pre-intervention period. In addition, the control time series must not themselves have been affected by the intervention. For details, see [Brodersen et al. \(2014\)](#).

A simple example


The figure below shows an application of the R package. Based on the observed data before the intervention (black) and a control time series (not shown), the model has computed what would have happened after the intervention at time point 70 in the absence of the intervention (blue).

The difference between the actual observed data and the prediction during the post-intervention period is an estimate of the causal effect of the intervention. The first panel shows the observed and predicted response on the original scale. The second panel shows the difference between the two, i.e., the causal effect for each point in time. The third panel shows the individual causal effects added up in time.

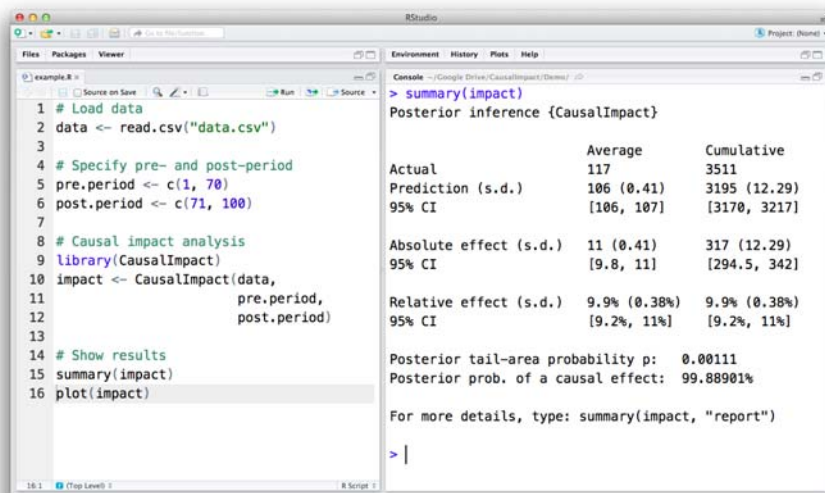

[Labels](#)
[Archive](#)
[Feed](#)


googleOSPO

YouTube 3K



The script used to create the above figure is shown in the left part of the window below. Using package defaults means our analysis boils down to just a single line of code: a call to the function `CausalImpact()` in line 10. The right-hand side of the window shows the resulting numeric output. For details on how to customize the model, see the [documentation](#).



How to get started

The best place to start is the [package documentation](#). The package is hosted on [Github](#) and can be installed using:

```

install.packages("devtools")
library(devtools)
devtools::install_github("google/CausalImpact")
library(CausalImpact)
    
```

By Kay H. Brodersen, Google

Labels: [analytics](#), [r](#), [statistics](#), [time series](#)

18 comments:

Seriously... said...

Thank you for sharing this. :]

September 12, 2014 at 7:54 PM

dan said...

Hi guys, there will be a verion of CausalImpact for the version 2.14.3?

September 14, 2014 at 9:23 AM

Mehmet Kaplan said...

Very nice...

I wonder if multi causals exist and the impacts have lag (or say phase difference) then still the code is useful or not.

If you think it is useful, can you please share a sample?

September 14, 2014 at 2:23 PM

Sergei said...

instal.package("Boom") is broken

<http://stackoverflow.com/questions/24355566/r-cmd-install-make-error-1-error-compilation-failed-for-package>

September 15, 2014 at 5:49 AM

Nick said...

Anybody having issue downloading this package under R version 3.1.1 ? I am currently using that version but unable to install it...

September 16, 2014 at 6:00 AM

Kay Brodersen said...

Hi Dan, thanks for reaching out. CausalImpact requires R >= 3.0.2. Kay

September 18, 2014 at 8:04 AM

Kay Brodersen said...

Hi Mehmet, the package is primarily designed for single interventions. Causal attribution in the presence of multiple interventions is possible but does introduce additional subtleties. See the Google group for a recent discussion (<https://groups.google.com/d/forum/causalimpact>). Kay

September 18, 2014 at 8:08 AM

Kay Brodersen said...

Hi Nick, the best way forward would be to post a reproducible example including details about your platform to the Google group. Kay

September 18, 2014 at 8:12 AM

Diogo said...

Thanks for sharing! I was wondering if anyone has test this method against more traditional randomized control trials for example?

September 18, 2014 at 11:49 AM

Kay Brodersen said...

Hi Diogo, yes, this is being done a lot - it's a great way of validating the approach for a given environment. For example, you could run CausalImpact on a time series that was acquired under RCT conditions, then compare the resulting counterfactual estimates to the original control group. See Section 4 in the paper for an example. Kay

September 19, 2014 at 6:28 AM

Dr Clements said...

Sounds good, tho' limited (in only having one intervention). But more importantly, what's needed is stats for non R-Geeks - something wherein an Excel set of data can be poured in and stat significant relationships between independent and dependent variables sussed out

September 23, 2014 at 10:19 AM

Dave Clarke said...

Hi, I am receiving the following error when I try to install with R v3.1.1 - 'ERROR: dependency 'bsts' is not available for package 'CausalImpact' Can someone help?

October 6, 2014 at 3:26 AM

www.phillipburger.net said...

Successful install on:

R version 3.0.2 (2013-09-25)

Platform: x86_64-unknown-linux-gnu (64-bit)

Will be running some data in next few days. Will post again if I run into info that can help answer any of the above questions.

October 15, 2014 at 11:34 AM

Giorgio Ballardin said...

Have you repeated your empirical analysis in "Analysis 3: Absence of an effect on the controls" in section 4 over a large number of time series that had no intervention? The single data point you quote in your paper under this analysis is interesting, but I'd be curious to see the distribution of the model results over, say, 10,000 of such but different comparisons.

October 17, 2014 at 12:49 PM

andre violante said...

We're doing some advertising tests with test and control groups very similar to the example in your publication except we're doing state tests and not DMA. I just have a couple questions as I'm reporting the analysis:

1. Is it okay to re-run the model every several days?
 - a. I noticed the historical point pred and point effect changing each iteration.
 - b. Pre.period and post.period remain the same except post is extended like so: `(post.period <- c(366, nrow(dispatch_dat)))`
2. Is there a minimum or maximum pre.period amount of data for training? What is best practice?

December 11, 2014 at 4:29 PM

Kay Brodersen said...

Re: Giorgio Ballardin

This is a great idea. We do null analyses routinely in every study, so I'd encourage you to try these at the scale you suggested.

December 17, 2014 at 12:45 PM

Kay Brodersen said...

Re: andre violante

I'd suggest posting this question in the CausalImpact Google group where we'll have more room to go into details.

<https://groups.google.com/forum/#!forum/CausalImpact>

December 17, 2014 at 12:47 PM

andre violante said...

Is there a way to set a model constraint that the prediction must be > 0 since the data is clicks, visits, revenue, etc. and a value < 0 would be impossible?

January 15, 2015 at 11:30 AM

[Post a Comment](#)

Links to this post

[Create a Link](#)

[Newer Post](#)

[Home](#)

[Older Post](#)



Company-wide

[Official Google Blog](#)

[Public Policy Blog](#)

[Student Blog](#)

Products

[Android Blog](#)

[Chrome Blog](#)

[Lat Long Blog](#)

Developers

[Developers Blog](#)

[Ads Developer Blog](#)

[Android Developers Blog](#)